

UNITED STATES PATENT APPLICATION

of

Timothy Ray Locascio

James DellaMorte

Rajnish Jain

John DeCesare

and

Gary Baldino

for a

**MEDIA RESOURCE CARD WITH PROGRAMMABLE CACHING FOR
CONVERGED SERVICES PLATFORM**

MEDIA RESOURCE CARD WITH PROGRAMMABLE CACHING FOR CONVERGED SERVICES PLATFORM

BACKGROUND OF THE INVENTION

Field of the Invention

5 The present inventions relates generally to the field of telecommunications and, more specifically, to a media resource card which may be used, in conjunction with a converged services platform, to provide a wide variety of enhanced communications services.

Background Information

10 A continuing trend of convergence between the circuit-switched public switched telephone network (PSTN) and packet-switched networks, such as Internet Protocol (IP) networks, has created demand for systems capable of interfacing with both types of networks while supporting a large number of diverse applications. Such applications include toll-free calling, wireless/wireline prepaid calling, directory assistance, voicemail,
15 call centers, conferencing and many others. In order to support such diverse applications, it is generally necessary to provide a capability for performing DTMF tone generation/detection, voice recorded announcements and similar communications services. Such basic capability, conventionally packaged as a "card" or printed circuit board which interfaces with an industry standard bus, has been commercially available for some time
20 from a variety of vendors including Intel Corporation and NMS Communications.

 However, marketplace demands indicate that greater scalability, density, smaller packaging and greater performance/cost ratios are needed in order for telecommunication carriers and other service providers to operate profitably and to be in position to rapidly deploy new applications which represent new revenue sources. As a result, there is need

for greater flexibility and programmability in delivering the communications services which enable the applications.

SUMMARY OF THE INVENTION

In brief summary, the present invention provides a media resource card which, when integrated into a converged services platform, is operable to perform a variety of enhanced telecommunication services. Such services include tone detection/tone generation, conferencing, recording and playback, and numerous others.

In a preferred embodiment, the media resource card includes a plurality of digital signal processors (DSPs), each of which has an associated cache memory, and a communications interface to the main microprocessor. The main microprocessor supports a network interface through which the media resource card may communicate with a networked media (file) server. In accordance with one aspect of the present invention, a file may be retrieved from the media server or other source and cached in one or more cache memories. The cached file preferably has at least two programmable attributes associated with it: a persistence level and an expiration time. The persistence level attribute may be used to specify how readily or not a particular file may be deleted from the cache. The expiration time attribute may be used to specify a time period after which a particular file may be deleted from the cache.

In a preferred embodiment, both attributes are programmable by way of an application program. That is, an application program may set one or both attributes on a per file basis, thus permitting an application developer to optimize cache performance for a given application.

BRIEF DESCRIPTION OF THE DRAWINGS

The invention description below refers to the accompanying drawings, of which: Fig. 1 is a block diagram of a converged services platform which includes a media resource card constructed in accordance with a preferred embodiment of the present invention; and

Fig. 2 is a block diagram of the media resource card shown in Fig. 1.

DETAILED DESCRIPTION OF AN ILLUSTRATIVE EMBODIMENT

Fig. 1 shows a converged services platform 2 which is controlled by an application program (not shown) running on a host computer 4. Hardware and software which
5 may be used to implement a converged services platform of the type shown are available from Excel Switching Corporation of Hyannis, Massachusetts. With the inclusion of appropriate cards discussed below, platform 2 is capable of interfacing with both the PSTN 6 and an IP network 8.

As shown, platform 2 includes redundant switching buses 10a and 10b. Redundant CPU cards 12a, 12b are connected to buses 10a, an HDLC bus 14, and to host 4 by
10 way of input/output (I/O) cards 16a, 16b. Depending upon the requirements of a particular application, various combinations of the following "line" cards, each of which supports a particular digital telecommunications protocol, may be included within platform 2: T1 card 16; E1 card 18, DS3 cards 20a, 20b. For redundancy, a standby card 22,
15 which is a duplicate of one of the other line cards, may also be included. All such line cards have an associated I/O card, denoted collectively by reference number 26, which serves as an interface to PSTN 6.

If a particular application requires ISDN signaling, ISDN PRI cards 24a, 24b may be included within platform 2. Similarly, if a particular application requires SS7 signaling, SS7 cards 28a, 28b may be included. If a particular application requires IP media
20 transport, platform 2 may include IPN cards 30a, 30b or an IPS card 32. Also, if a particular application requires media services, a media resource card 38 may be included. Cards 30, 32 and 38 each have an associated I/O card, denoted collectively by reference number 40. Power cards 34a, 34b and cooling fans 36a, 36b are also present within platform 2.
25

Referring now to Fig. 2, a detailed block diagram of the media resources card 38 is shown. A DSP module 44a includes four DSPs 46a-46d, each of which has an associated cache memory 48a-48d, respectively. Each of DSPs 46a-46d is preferably a Texas

Instruments TMS320C6414, but any of a number of other commercially available DSPs could readily be adapted for use with the present invention.

DSPs 46a-46d are connected to a time division multiplex (TDM) switch 50. Each of DSPs 46a-46d is capable of receiving and transmitting up to 512 DS0s (time-slots) from TDM switch 50. If desired, a second DSP module 44b, containing an additional four DSPs (not shown) may be added to media resources card.

Overall control of the operation of media resources card 38 is carried out by a CPU 56, which is preferably a Motorola PowerQUICC 8260 processor, having an associated cache memory 58 and a local memory 60. CPU 56 also has an associated Ethernet network interface through which the CPU may communicate with a Network File System (NFS) file server 42. Any of a number of other commercially available microprocessors or other file system protocols, including CIFS or TFTP, could readily be adapted for use with the present invention.

Cache memories 48a-48d and 58 are preferably used to cache voice recorded announcements or other recordings which are intended for playback to callers (not shown) as directed by an application running on host computer 4 (Fig. 1). Large quantities of such voice recorded announcements may be stored on file server 42, in accordance with NFS, CIFS or another desired file system protocol, retrieved by CPU 56 and then cached at the time of initial playback.

We will now describe the operation of a preferred embodiment of a caching algorithm, preferably executed by CPU 56, which may be considered an adjunct to a conventional Least Recently Used (LRU) algorithm (*i.e.*, a file which is least recently used is normally the first candidate to be deleted from the cache memory). We shall consider as a first example the case of a "cache miss" in which a file needed for playback is not present in any of the cache memories 48a-48d or 58. First, CPU 56 must retrieve (read) the file from file server 42. Assuming that the file does not exceed a predetermined size, CPU 56 will cache the retrieved file in both cache memory 58 and the one of cache memories 46a-46d which is associated with the DSP that CPU 56 has assigned to initially playback the file. The assignment of DSPs by CPU 56 may be done in any of a number of ways including a "round robin" approach.

An advantage of caching the file in the CPU cache memory 58 is that, in the event the same file is needed for playback in the future, but playback is to be performed by a different DSP than the one that performed the initial playback, the file may be retrieved from cache memory 58 (thus avoiding another file server access) and cached in the appropriate DSP cache memory.

Contemporaneously with beginning playback of the file, either of two (or both) attributes are assigned to the file by CPU 56 in response to messages received the application running on the host computer 4. The first such attribute, referred to as a persistence level, is a value which specifies how readily or not the associated cached file may be deleted from one or more of said cache memories. In a preferred embodiment, one of four possible persistence levels is specified: none; soft; firm; and hard. A persistence level of "none" means that the cached file is transient and may be freely deleted to make room for other files. A "soft" persistence level means that the cached file should remain in the cache memory until the LRU algorithm deletes it. A "firm" persistence level means that the cached file should remain in the cache memory until all of the "soft" persistence level files have been removed by the LRU algorithm. A "hard" persistence level means that the cached file should not be deleted from the cache memory.

As noted above, persistence levels are preferably specified by an application program and are thus under the control of the application developer who, given his or her knowledge of the application's expected performance, can choose persistence levels to optimize cache performance. One exception to this arrangement is the case of a cached recording (*e.g.*, a voicemail message or a directory assistance call) made by the media resources card 38. Because a cached recording is "one of a kind" and is initially not stored anywhere except in one of the cache memories 48a-48d, CPU 56 will assign a "hard" persistence level to such a file to ensure that it is not deleted from the cache memory before being stored elsewhere or there is certainty it is no longer needed.

The second such attribute is a timer expiration, which simply represents a time period after which a cached file may be deleted. If set, the timer expiration attribute may override the persistence level attribute.

We will now consider a second example in which we have a “cache hit” meaning that a file that is needed for playback is present in at least one of cache memories 48a-48d and 58. In this example, CPU 56 will assign one of DSPs 46a-46d to perform the playback. If the cache memory associated with the assigned DSP contains the needed file, the
5 assigned DSP may proceed with playback. If the needed file is contained in the cache memory associated with a different DSP or cache memory 58, then file will be cached to the assigned DSP’s cache memory contemporaneously with playback.

What is claimed is: